

ANISH KUMAR

ML Engineer | AI/ML Final-Year Student

[linkedin.com/in/anishkumar-ai](https://www.linkedin.com/in/anishkumar-ai) | github.com/anishk-neuroforge | chief-anishkumar-ai.me | anish00kumar11@gmail.com | +91 7976242904

SUMMARY

Final-year BTech AI/ML student building production-oriented machine learning systems across the full lifecycle — data pipelines, model training, inference optimization, and API serving. Experience shipping end-to-end ML systems (RAG, NLP, CV, automation agents) with FastAPI, Docker, and modern MLOps tooling. Solved 100+ DSA problems on LeetCode/GeeksforGeeks. Seeking full-time ML Engineer roles.

SKILLS

- Languages: Python, C, C++, Java
- ML/DL Frameworks: PyTorch, TensorFlow, OpenCV, Hugging Face Transformers, LangChain
- Inference & Serving: FastAPI, ONNX, REST APIs, Postman
- MLOps & Data: Docker, MLflow, FAISS, MySQL, AWS
- Tools: Git, GitHub, Maven, DSA

EXPERIENCE

AI Training & Evaluation Specialist — Outlier AI2024 – 2025

Freelance / Part-time

- Performed RLHF-style evaluation of 500+ LLM responses for safety, reasoning quality, and code correctness.
- Structured feedback and rubric-based scoring improved prompt acceptance rate by 18%.

PROJECTS

Document Q&A — RAG Pipeline [DOC](#)

- Built an end-to-end RAG pipeline for multi-document Q&A using PyMuPDF extraction, tiktoken chunking, sentence-transformer embeddings, and FAISS vector search, generation via the Groq API.
- Improved answer quality from a 0.41 BLEU keyword-search baseline to 0.81 BLEU on a 1,000+ document PDF corpus, at ~1.2s end-to-end latency.

MedReport AI — Medical Report Analyzer

- Designed a 4-stage pipeline that turns raw lab-report PDFs/scans into plain-language explanations: PaddleOCR extraction, regex-primary entity matching (with an opt-in BioBERT NER pass for non-templated reports), anomaly scoring against a 50+ lab-test reference database, and LLM-based (Gemini API) explanation generation.
- Prioritized transparent handling of ML-component limitations over inflated accuracy claims; prototype achieves ~95% extraction accuracy across varied report layouts at sub-2s latency.

Predictive OS Resource Allocation [OS](#)

- Built a GUI-based OS resource allocator that predicts process behavior to optimize CPU and memory utilization.
- Designed dynamic scheduling logic anticipating future resource needs, with a real-time visualization dashboard for process queues, CPU usage, and memory distribution.

JAS — Natural-Language Desktop Automation Agent [JAS](#)

- Built an LLaMA 3.1-powered agent that translates natural-language instructions into desktop and browser automation actions.
- Re-architected the browser-automation subsystem (Browser V2) with Playwright, DOM-first observation, a semantic locator waterfall, and deterministic DOM diffing to replace a fragile vision-based control loop.

TRAININGS & CERTIFICATIONS

RHCSA — Mastering Linux System Administration | Red Hat, Inc. Jun 2025 – Jul 2025

- Hands-on Linux system configuration, user management, security administration, shell scripting, and process automation.

ACHIEVEMENTS

- Solved 100+ Data Structures & Algorithms problems across LeetCode and GeeksforGeeks (LPU, Nov 2025).

EDUCATION

Lovely Professional University — B.Tech, CSE (AI/ML) • CGPA: 6.29

Since Aug 2023

Sri Guru Harkrishan Sr. Sec. Public School, Chandigarh — Intermediate (PCM) • 70.20% *Apr 2020 – Mar 2022*